

Information Power Grid

*Distributed High-Performance Computing and
Large-Scale Data Management for Science and Engineering*



William E. Johnston,
Dennis Gannon, and Bill Nitzberg
Numerical Aerospace Simulation
Division

William J. Feiereisen, Division Chief,
William Thigpen, Engineering Branch Chief,
Alex C. Woo Research Branch Chief
<http://www.nas.nasa.gov/IPG>

What is the Information Power Grid?: Prototype (cont.)

The CX testbed will address aggregated large system computing and data management:

- CPU resource reservation
- network resource reservation
- standardized access to high capacity storage
- a stakeholder-oriented, distributed resource management approach that addresses *global use but local control*
- operational tools and procedures and user support
- allocation, accounting, and auditing
- strong user identification and access control
- model for CoSMO [23]

Prototype and Testbed Environments

The IPG *prototype production* environment (the “CX” IPG testbed) targets an *operational Grid environment incorporating major computing and data resources at multiple NASA sites* in order to provide an infrastructure capable of routinely addressing larger scale, more diverse, and more transient problems than is possible today.

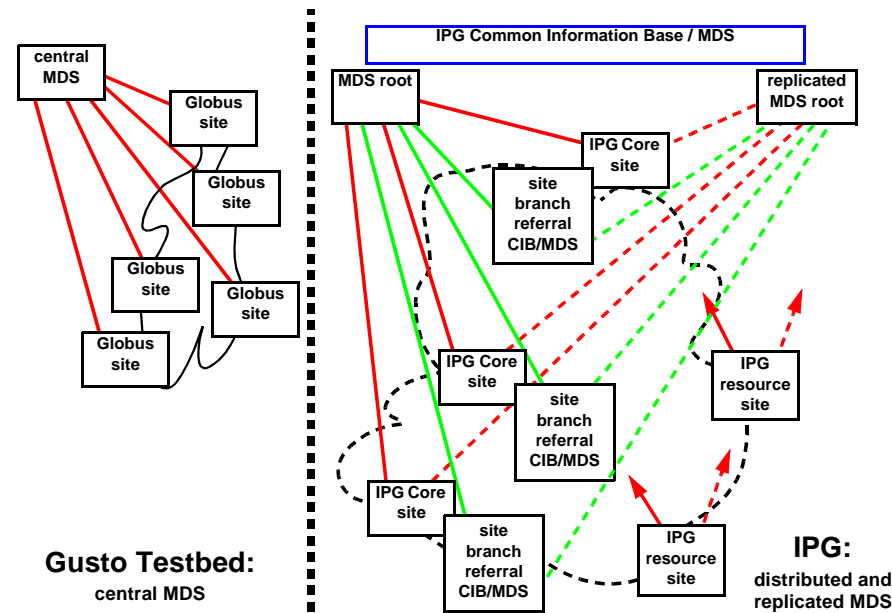
What is the Information Power Grid?: Prototype (cont.)

IPG testbeds will also provide R&D platforms:

- *Dev* testbed: IPG system software development
- *DX* testbed: high data-rate distributed resources
 - high-speed end-to-end
 - high data-rate services including data archives and instrument systems
 - test applications
- *SX* testbed: security and infrastructure protection

State of IPG: Baseline Operational System (cont.)

- .. Globus providing the Grid common services
- .. Programming and program execution support
 - Grid MPI (via the Nexus communications library)
 - CORBA integrated with Globus
 - global job queue management
 - high throughput job manager
 - Condor [24] (“cycle stealing” computing)
- .. A stable and supported operational environment
- .. Several “benchmark” applications operating across IPG (multi-grid CFD code, parameter study)
- .. Multi-Grid operation (applications operating across IPG and NCSA)



Gusto Testbed:
central MDS

Migration of Common Information Base from R&D to Prototype Production

How Are We Accomplishing the Baseline?

1) Persistent operational environment that encompasses significant resources

- .. “IPG Prototype Startup Tasks (target: 6/99)
 - Globus deployed across Ames, GRC, and LaRC (Task 1.0)
 - IPG common grid information base (Task 2.0)
 - IPG X.509 Certification Authority and certificate server (Task 3.0)
 - Global queuing and user-level queue management capability on top of Globus (Task 4.0)

Accomplishing the Baseline - Operational Environment

- Computing resources for the initial IPG multi-center testbed (CX) (Task 5.0)
- Networking for the IPG Testbed: Ames, GRC, LaRC (Task 6.0)
- IPG Access for Archival and Published Data: SDSC’s Metadata Catalogue (MCAT) and the Storage Resource Broker (SRB) (Task 7.0)
- Heterogeneity in the IPG testbed: Condor (Task 8.0)
- Heterogeneity in the IPG testbed: High performance clusters (Task 9.0)

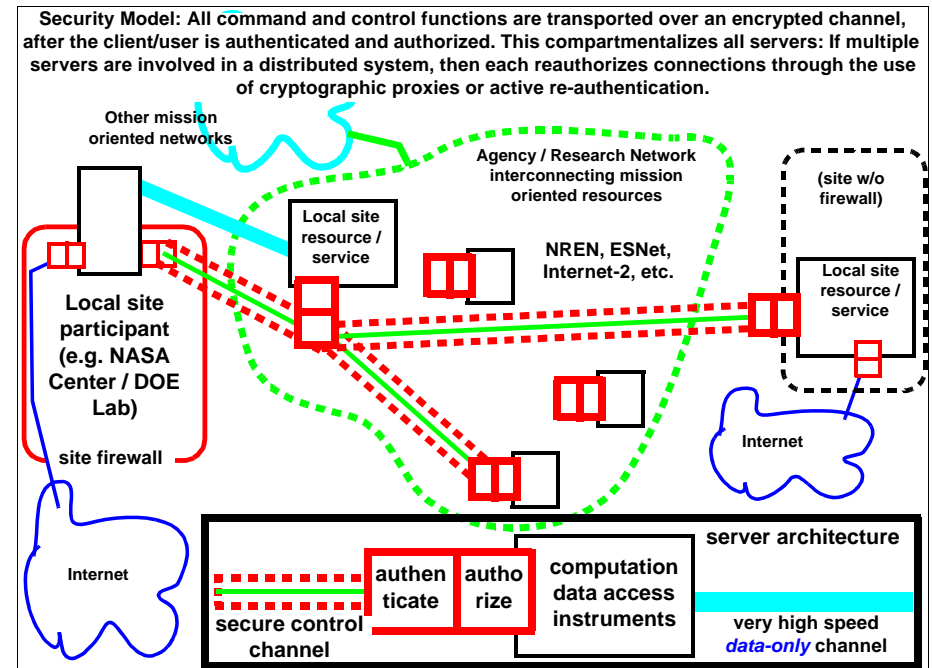
Accomplishing the Baseline - Operational Environment

IPG Operational Tasks (needed for v.1.0 of IPG baseline - 9/99)

- Security (Task 10.0)
- IPG Information Base / MDS database maintenance (Task 11.0)
- IPG/Globus system administration (Task 12.0)
- Automatic monitoring of IPG components (Task 13.0)
- Trouble ticket model (Task 14.0)
- Condor support (Task 15.0)
- CORBA support (Task 16.0)
- Legion support (Task 17.0)
- Documentation (Task 18.0)

Accomplishing the Baseline - Operational Environment

- User services (Task 19.0)
- Account management (automated generation and maintenance mechanisms) (Task 20.0)
- Globus with multiple MDS and PKI (Task 21.0)
- Allocation Management and Accounting (Task 22.0)
- System testing: Verification suites, benchmarks, and reliability/sensitivity analysis for IPG (both static and dynamic) (Task 23.0)



Accomplishing the Baseline - Operational Environment

IPG Functionality Tasks (near-term - 9/99)

- CORBA in the IPG environment. (Task 24.0)
- Integration of Legion (Task 25.0)
- CPU resource reservation (Task 26.0)
- High Throughput Computing (Task 27.0)
- Programming Services (Task 28.0)
- Distributed debugging (Task 28.1)
- Grid enabled visualization (Task 28.2)

Accomplishing the Baseline - Operational Environment

- .. IPG Functionality Tasks (mid-term - 9/00)
 - Network bandwidth reservation (Task 29.0)
- .. Characteristic Applications - 9/99
 - OVERFLOW port & tune (Task 31.0)
 - NPSS port & tune (Task 32.0)
 - Parameter study
 - Heterogeneous testbed application: Condor (Task 34.0)
 - Heterogeneous testbed application: High performance clusters (Task 35.0)

How Are We Accomplishing the Baseline? - Service Model (cont.)

Implementing this service model requires two things:

- **tools for production support** of integrated collections of widely distributed resources must be identified, built, and provided to the systems and operations staffs
- **new organizational structures** must be evolved that account for the fact that operating Grids is different than operating traditional supercomputer centers, and management and operation of this new shared responsibility service delivery environment must be explicitly addressed

How Are We Accomplishing the Baseline? (cont.)

2) New service delivery / operational models

Grids such as IPG, effectively define a new business model for operational organizations that will deliver large-scale computing and data resources *in ways that allow them to be integrated with other widely distributed resources controlled*, e.g., by the user community.

How Are We Accomplishing the Baseline? (cont.)

What can the commercial sector provide that will help?

- .. Participate in the Grid Forum to learn about the issues that are facing Grid developers
- .. Batch schedulers with well defined APIs supporting advance reservation
- .. MPI highly tuned for the vendor's environment AND that will interoperate with other vendor's implementations
- .. Support for highly optimized SSL

- .. Highly tuned TCP that supports window scaling
- .. Open and optimized parallel ftp with TCP options and PKI/X.509 authentication
- .. Support for SSH with PKI/X.509
- .. LDPA servers that really work and support high update rates and referral
- .. PKI interoperability - participate in Open PKI group (<http://www.pkiforum.org>)
- .. High performance, multi-threaded CORBA ORBs on the supercomputing platforms
- .. Secure CORBA (SSL + OMG object call-out)

How Are We Accomplishing the Baseline? (cont.)

3) New functionality will arise from:

- .. A *detailed and on-going examination of requirements* generated by several NASA application communities, both in terms of specific capabilities identified by the applications community
- .. Continual analysis of the requirements and desired operating environments by computer scientists
- .. Research and development to address the identified needs

- .. Tertiary storage systems that support high performance to WANs (multi-gigabit data streams)

The Two Year IPG Goals

- .. An *operational and persistent, “large-scale” prototype-production Information Power Grid providing access to computing, data, and instrument resources at NASA Centers around the country*
- .. A well developed R&D program addressing open issues in Grid computing and data management: i.e., defined tasks and progress in the areas listed under “*Research and development in support of long term goals*,” above.

- .. Applications that cannot be done today:
 - *computational grid-independent rotocraft reference simulation*
 - *multi-component, multi-disciplinary turbo-machine simulation*
 - *configurable problem solving environment for large parameter space studies (e.g., RLV [26])*
 - *on-line instrument (Darwin wind tunnel system) coupled to scheduled computational and data resources*
- .. Transfer of Grid technology to the production environment (CoSMO)

What Will IPG Facilitate? (cont.)

- .. *Routine use of wide area, data-intensive applications*
 - remote access to high data-rate real-time data sources / instruments and very large datasets
- .. *Building and managing just-in-time, dynamic, and collaborative systems*
 - coordinated work by dispersed groups based, e.g., on central design databases (e.g. airframe or turbomachinery geometry and performance)
 - data and simulation based crisis response

What Will IPG Facilitate?

- .. Building just-in-time, large-scale systems/applications to support scientific and engineering *computing and data oriented activities* that are not steady state, i.e. those that may require, or have to make use of, a different resource mix for every different problem – e.g.:
 - *coupled, multidisciplinary simulations*
 - *large-scale simulations*
 - *coordinated use of many dispersed data archives*
 - *data analysis based control of on-line instruments*
 - *coupling of remote, on-line instruments to large-scale computation simulation*

“Top Twenty” R&D topics for IPG

These topics address functionality that is of primary importance for some aspect of IPG in the first two years in order to satisfy an explicit requirement (either from user community or from operations/systems) or an implicit requirement (clearly needed to support our Grid model). This is not an exhaustive list of such topics because it excludes topics covered in the Engineering WG (functionality that is needed in the first year) and it does not cover topics that are being worked on by IPG funded partners (e.g. the NSF PACIs).

“Top Twenty” R&D topics for IPG (cont.)

Tools for Converting and “Wrapping” Legacy Codes for Operation in Grids

- Most existing simulations are codes that have evolved over many years, together with the understanding of the underlying physics of the problem. As such, few of these codes are “modern” in design: single function code units, well documented and recoverable exceptions, re-entrant, etc. Yet these codes must be usable in modern distributed computing environments in order to make use of the available computing and storage resources.

Toolkits for Constructing Problem Solving Environments / Workbenches

- Provide the tools and techniques for integrating Grid applications and services into the desktop environment in ways that can express and facilitate the human work process. This includes: General graphical user interface components for providing the application interface; Scripting / PSE program preservation and editing; Interfaces to generalized workflow management services; Tools for managing high throughput jobs; Global shell” functionality; Rule-based workflows driven from published/subscribed events; Resource discovery and brokering services; Data cataloging and access functions; Access control functions.

Application-Level Fault Management

High-Speed Application Access to Data Files: Staging, Caching, Location Management, and Remote I/O

- Data location management is a very important facility in the Grid, and will probably be one of the services that most determines whether the Grid can truly be successful in providing widely distributed access to computing, data, instrument, and human resources.

Establishing the Execution Environment

- Tools are needed to set up the environment in which code must actually execute: binaries have to be staged, runtime libraries have to be instantiated, file system environment must be established, environment variables and other state needed must be established. This situation presents both the problem of how to describe what is needed and then to establish all of the required environment on the target system.

Application Performance Monitoring

- There are several aspects to monitoring. Characterizing the operating state and predicting future state of the execution environment is important for supporting execution management such as fault management and for adaptation and performance analysis of running algorithms. Analytical monitoring is important for both algorithm characterization and troubleshooting. Characterization and prediction is provided, e.g., by the Network Weather Service [86] and characterization analysis is typified by the precision time event tagging for dispersed, multi-component performance analysis provided by Netlogger [85]. Further, monitoring is likely to be important for functionality such as generalized

- Beyond the process-existence oriented monitoring and management approach of the previous section, in multi-element production software systems (such as DAO’s data acquisition, validation, and modeling for daily weather forecasting) application specific faults must also be dealt with. That is, mechanisms are required for managing generalized “faults” in all aspects of the application environment. The concept of generalized faults is intended to include things like application problems such as failure to achieve numerical convergence within a specified CPU time utilization, non-availability of resources or data, etc. These and other types of faults must be detected and acted upon to provide for recovery (continued operation of the application). The definition of these faults and the recovery procedures will be application specific, so what is needed is a framework that provides tools for fault definition and detection, followed by knowledge/rule based recovery.

Global Shells and Generalized Workflow Management

- The Grid workflow management mechanism must provide for description and subsequent control of the ordered steps and events that represent a “job”. An approach analogous to shells in our local system environments might have Unix shell-like semantics, but global reach and naming. A more general approach is represented by a rule-based execution management system driven from published/subscribed global events (where the “events” represent process completion, file or other state creation, instrument turn-on, etc.).

Global Event Publish/Subscribe Systems

auditing e.g., data file history and control flow tracking in distributed, multi-process simulations. Work is needed on what should be done, and how existing approaches could be integrated into IPG.

Global Queuing and Execution Queue Management

- There are a wealth of queue management related issues and services that must be addressed in order to make Grids credible production environments and to provide better reliability and resource utilization than is available today. While some of these issues are arguably the purview of resource management, global queue management is resource independent, and is therefore included in execution management.

Support for Multiple Programming Paradigms in Multi-platform, Heterogeneous Computing Environments: Globus, CORBA, Condor, Java, and Legion

- The most common programming paradigms identified thus far in the NASA environment are MPL/Fortran, MPI/Fortran, and Java. CORBA is used to wrap codes so that standardized interfaces may be defined. In addition to Globus [55], Java, Condor [37], and Legion [78] are actively used programming paradigms that will be supported in IPG. In the case of CORBA support means ensuring co-existence with IPG services and investigating and implementing cooperation with Grid services. In the case of Java support means interfaces to IPG services. Condor is already integrated. In the case of Legion, support in the IPG environment means that Legion will make use of the Grid Common Services such as the Common Information Base, job submission and management, etc.

Network Bandwidth Advance Reservation Scheduling: Differentiated IP Services

- Reservable, high bandwidth data flows are essential to support transient network uses such as on-line scientific instruments that produce data only during experiments, specific data analysis exercises, or for large distributed computational runs when co-scheduling computing and communication resources will be necessary to keep the processors operating at a reasonable efficiency. These are scenarios in which an ensemble of data sources, and computing and storage resources must be connected and operated as a system — that is, all resources must be available at the same time as must the communications services that provide the interconnections to form useful distributed systems and to provide the “agility” to support transience (systems built on-demand for limited periods). Such reservation capability must be provided as a standard IPG service.

CPU Advance Reservation Scheduling

- The rationale for advance CPU reservation is the same as advance network bandwidth reservation: This is an essential capability to support building aggregated and just-in-time systems. This capability must be a standard IPG service.

High Throughput Computing / Parameter Study Support

- High-throughput computing / parameter study assigns independent tasks to multiple independent computing resources and schedules and manages them so that a large number of similar jobs execute in parallel. Sometimes a goal is to use otherwise unused computer cycles. Aerospace system design often requires searches through

throughput jobs; Global shell” functionality; Rule-based workflows driven from published/subscribed events; Resource discovery and brokering services; Data cataloging and access functions; Access control functions.

Access Control

- Access control is essential in order to address several issues in Grid environments. One obvious issue is that owners of confidential or proprietary data must have strong and easily managed mechanisms that ensure only authorized users gain access to that data. From the point of view of the Grid itself, management of stakeholder rights (use-conditions) is an essential aspect of having individuals and institutions make their resources available as part of any Grid environment. Many resources may only be utilized by people/communities that satisfy certain criteria as set by funding agencies, etc., as a matter of policy. Tools, techniques, and infrastructure must be built into IPG from the start in order to address these issues.

Distributed Collaboration Tools and Techniques

- Toolkits supporting the construction of PSEs must provide the mechanisms for integrating computer mediated, distributed human collaboration into desktop problem solving environments. E.g. interface sharing, graphical user interface components that map to applications and Grid services, access control, a representation of the human work process that maps onto the workflow management mechanism, etc.

parameter space. In such cases, since each parameter is independent, individual or groups of parameters can be assigned to separate computers and then processed independently.

Data Cataloging and Publication Services

- Services are needed to provide:
 - automatic characterization of datasets and “publication” of the characterization where, and in a form, that it may be located
 - description of data (generate metadata) and data formats
 - management of dataset use conditions and access control

Grid Distributed Debugging

- There currently are no commercial or experimental tools that can be directly incorporated into a Grid programming environment that can provide the needed level of support for distributed system debugging. Consequently, research and development is needed to adapt the current generation of parallel debuggers for dynamic and adaptive Grid applications.

Remote Visualization

- Provide the tools and techniques for integrating Grid applications and services into the desktop environment in ways that can express and facilitate the human work process. This includes: General graphical user interface components for providing the application interface; Scripting / PSE program preservation and editing; Interfaces to generalized workflow management services; Tools for managing high

High-Speed Application Access to Data Files: Staging, Caching, Location Management, and Remote I/O

- Data location management is a very important facility in the Grid, and will probably be one of the services that most determines whether the Grid can truly be successful in providing widely distributed access to computing, data, instrument, and human resources.

Partners

NASA partners:

- Glenn Research Center
- Langley Research Center

Funded collaborations:

- NSF PACIs via NASA / NSF MOU
(Alliance/NCSA and NPACI/SDSC)
- U. Ill., U. Chicago, U. Wisc., Argonne National Lab.,
UCSD, USC/ISI, U. Tenn., U. Va.

Others:

- Lawrence Berkeley National Lab (network QoS)